"Response, Recovery and Reducing Our Vulnerability to Cyber Attacks: Lessons Learned and Implications for the Department of Homeland Security"

Statement of Alan Paller
Director of Research
The SANS Institute

to the United States House of Representatives
Subcommittee on Cybersecurity, Science, and Research & Development
Select Committee on Homeland Security

June 25, 2003

Chairman Thornberry, Congresswoman Lofgren, distinguished Members of the Committee, I appreciate the opportunity to appear before you today. It is particularly gratifying to us in the cybersecurity field, Mr. Chairman, that a person with your foresight, vision and leadership in homeland security has decided to take on the challenges of cybersecurity. I am not sure whether my colleagues are aware that six months before the September 11, 2001 attack, you introduced a bill in the House of Representatives that called for consolidating the federal agencies responsible for protecting our homeland. You saw the threat clearly; you spoke eloquently of the technological dimension, but it took a major attack before others were able to share your vision. I am very hopeful that in the cybersecurity field progress can be made more quickly. With your leadership and that of Congresswoman Lofgren, who has been one of the most effective Members of Congress on high tech issues and whose district includes one of the largest concentration of computer companies and cyber security expertise anywhere in the world, Congress can help the Department of Homeland Security lead a rapid effort to reduce this nation's vulnerability to cyber attacks, turn the tide against cyber attackers, and increase our speed and effectiveness in responding to and recovering from the attacks that do succeed.

My name is Alan Paller and I am director of research at the SANS Institute. SANS is an educational institution. Last year, more than 14,000 system administrators and computer security professionals, from nearly every government agency and large commercial organization in the US and from 42 other countries, spent a week or more in SANS immersion training. They learned the details of attacks that will likely be launched against them, learned how to build and manage defenses for those attacks, and learned how to respond once an attack has occurred. SANS 38,000 alumni are on the front lines in the fight against cyber attacks. Once they have returned to work, we continue to support them and more than 120,000 of their coworkers with early warnings of new attacks, weekly summaries of new vulnerabilities and a research program that makes available more than 1,400 timely security research briefs.

In 2001, SANS created the Internet Storm Center, a powerful tool for detecting rising Internet threats. Storm Center uses advanced data correlation and visualization techniques to analyze data collected from more than 2,000 firewalls and intrusion detection systems in dozens of countries. Experienced analysts constantly monitor the Storm Center data feeds and search for anomalies. When a threat is detected, the team immediately begins an extensive investigation to gauge the

threat's severity and impact. Critical alerts are disseminated to the public via email and through the online press.

In my remarks today, I will share some of the successes and failures of the defensive community in responding to large cyber attacks, and I'll suggest ways that the lessons we learned might lead to effective initiatives for the Department of Homeland Security in improving response, recovery, and prevention.

Five months ago today, the Slammer worm attacked computers running Microsoft's widely used database management system. A worm, for those not steeped in the jargon of cyber security, is a malicious program that spreads from computer to computer without requiring users to take any action at all. Slammer represented a significant advance in attack technology. At its peak it was scanning 55,000,000 systems per second and that was 100 times as fast as Code Red scanned in July, 2001. Slammer infected 90% of the systems that were vulnerable in the first ten minutes of the attack and ultimately infected a total of 75,000 hosts. Slammer reminded the defensive community that we are engaged in an arms race with the attackers – one the attackers are likely to continue for many years. It did not contain a destructive payload; if it had thousands of organizations would have lost valuable data.

Slammer's high intensity scanning continued to wreak havoc for days. It surprised many people when it showed them that the computer systems that make up the nation's critical infrastructure for banking and transportation and emergency management - that some naively presumed to be somehow separate and isolated - are actually connected to the Internet and can be significantly affected by Internet attacks. For example, because of Slammer, Bank of America's ATM machines stopped dispensing money, Seattle's emergency 911 system stopped working, Continental Airlines had to cancel some flights because its electronic check-in system had problems, and Microsoft couldn't activate user licenses for Windows XP. Those were just a sample of the more high profile problems. Many other organizations were damaged by Slammer, but they managed to stay out of the press. The cult of secrecy that surrounds cyber attacks is part of the challenge we face in determining costs and in helping people recover.

On a more positive note, Slammer brought our focus back to two valuable lessons. The first, learned in the summer of 2001 as we responded to the Code Red worm:

1. Federal and private security specialists, working together, can create a synergy that doesn't appear to exist when they act separately.

Slammer did a lot of damage, but it did much less damage than it would have if government and private industry had not worked together to fight it. A team of private sector experts from large internet service providers (ISPs) discovered the worm when it started flooding their networks. Within minutes they contacted technical experts in government and CERT/CC (Computer Emergency Response Team Coordination Center), and those three groups joined forces to analyze the problem. They learned that the worm targeted a specific entry point on each computer, and that they could stop most of the damage it was doing by blocking traffic to that entry point. The ISPs reconfigured their networks to stop all network traffic destined for the worm's target entry point, and their customers – at least the ones that did not have their own infected systems – stopped feeling the pain.

For Slammer, early discovery, effective analysis and widespread notification led to immediate extensive filtering of the worm traffic, and that action protected many organizations from being overwhelmed. This worked so well on Slammer that one might well ask why we do not use the same approach on all large, automated attacks. The answer is that two barriers get in the way and both can be eased by Department of Homeland Security initiatives.

The first barrier is that the high speed filtering used for Slammer does not work for many other attacks. Slammer exploited a special path that could be blocked easily by existing network routers, without harming valid traffic. The Code Red worm, on the other hand, exploited the path universally used to request web pages. Anyone who blocked that path would stop all web traffic to their site. For an organization that uses the web for business purposes, blocking that path could inflict more damage than the worm could cause. To filter for Code Red and other worms that use popular paths, the network infrastructure used by large companies and ISPs needs to be upgraded so that it can selectively block malicious traffic. That type of high-speed, intelligent filtering is not yet widely available from the network equipment manufacturers. The Department of Homeland Security could help speed the availability of high speed filtering routers through research support and targeted procurement.

The second barrier is that the government and the rest of the defensive community cannot respond to attacks if they do not know that attacks are occurring. Slammer flooded huge numbers of systems, so it was easy to find. Attacks aimed at electric power grids or e-commerce sites or emergency response networks are not nearly as visible. Early warning for targeted attacks is possible only if the first victims choose to report the attacks rapidly. But just as people infected with communicable diseases are loathe to make spectacles of themselves, so victims of cyber attacks can see insufficient benefit in making their pain public even to government officials who promise not to tell others.

How can we increase prompt reporting on cyber attacks? Let's take a closer look at the medical analogy. People who become sick, even with a highly communicable disease, do not usually call the Center for Disease Control. But their doctors do make the call, and the doctors maintain the confidentiality of their patients' identities. In the cyber defense arena, consulting companies serve as doctors to help companies analyze cyber attacks and recover from them. This year, the Department of Homeland Security (DHS) is spending millions of dollars to create a Cyber Warning Information Network (CWIN) that connects organizations active in cyber defense so they can get early access to important information. To ensure the "doctors" report attacks to the DHS, the Department could require that organizations that want access to CWIN must commit to providing immediate notification to DHS whenever they or one of their clients is attacked, without naming the victim.

===

Slammer also reminded us of another significant lesson we learned in responding to Code Red and Slapper and many other worms:

2. A severe shortage of individuals with technical security skills combined with a lack of management focus on security issues, prevents many organizations from fully recovering

from attacks and improving their security. Better training is a partial solution, but joint action by government and industry to *standardize security configurations* and *automate patching* is already having a much larger impact.

Most attacks that do a lot of damage, like Slammer and Code Red, exploit vulnerabilities that are widely understood and for which remedies are known. Therefore it is surprising that two years after the Code Red worm swept through the Internet infecting vulnerable systems, 30,000 of those systems are still infected and still searching for other systems to infect. The problem is that many organizations that own computers have no one who understands how to secure those computers. When we find a Code Red infected system and ask why it hasn't been fixed, we usually hear that the system owner didn't know that it is attacking other systems and also that there is no one with security skills available to fix it.

Even large organizations are security-challenged. Slammer's victims included several huge security-sensitive organizations; Bank of America and Microsoft are examples. Their systems were flooded because vulnerable software had not been patched and because they had not configured their firewalls to block unwanted traffic from the Internet. It is unreasonable to blame the software users in this case, because Microsoft made installing this particular patch an arduous task, much more difficult than installing the underlying software in the first place. And most users and system administrators had never been told they should block the offending traffic at the firewall.

Training is part of the answer. Security-savvy system administrators are very effective at keeping their systems running smoothly while maintaining their defenses, and several large organizations are now requiring all system administrators to demonstrate their mastery of security as a prerequisite for getting control of the systems. However, most computers are not managed by system administrators. They are managed by busy people with other responsibilities. I do not believe it is fair or wise to expect that every graduate student or scientist or librarian who tries to install a workstation should become a security expert. And what about the grandparents and teenagers and all the other people who simply want their computers to work? We cannot ask them to develop and maintain the technical security skills needed to configure their systems safely and keep them secure.

A better solution is to remove the pain of security by centralizing and standardizing safe configuration and security patching. Large organizations can do that themselves, as the Department of Energy and others are demonstrating. But few other organizations have the time and talent. Only the companies that sell computers and software are positioned to make security configuration and patching inexpensive and effective.

Happily for all of us, vendors are beginning to recognize that security is a critical market need, and they are putting their development dollars to work to help their clients with security. Three weeks ago at a Federal Trade Commission workshop, Dell announced it would sell Windows 2000 systems configured in accordance with consensus security benchmarks, improving security and reducing the security burden for Dell customers. Similarly, Oracle and the Department of Energy are partnering to deliver safe configurations of Oracle software to all users at all Department of Energy laboratories and offices. Other Oracle users will benefit as Oracle makes the safer version available to the general public. Both of these efforts were facilitated by an extraordinary public-

private partnership involving the National Security Agency, the Defense Information Systems Agency, the Department of Energy, NIST, FedCIRC, SANS, and the Center for Internet Security (CIS). The CIS partnership has developed consensus benchmarks for safe configuration of many common operating systems and applications. Dell, for example, said that they would not have been able to create the new safer version of Windows 2000 without the work of the CIS partnership.

And automated patch delivery is maturing. For example, Red Hat delivers security updates for its software automatically as does Microsoft for some its Windows XP software.

It is common practice today for vendors to sell software and hardware with insecure configurations. Most users are not security experts and therefore are not aware of the configuration dangers, nor do they have the knowledge to find and apply appropriate security patches. All that means that millions of computers are at risk, and each of those vulnerable systems can be used by attackers to launch major denial of service attacks. With active leadership by the vendors and the federal government, worms and automated attacks will be denied easy access to all these systems. So what can the Department of Homeland Security do to accelerate this beneficial trend? DHS can require its vendors to deliver safe systems out of the box and ensure that patches are delivered automatically. As other federal agencies and companies follow the DHS lead, the market will reward vendors that take security burdens off their customers' shoulders.

===

In your letter of invitation, you also asked me to address the challenges in estimating the damage done by cyber attacks and the strengths and weaknesses of simulations and exercises for cyber security. I'll answer both briefly because the general knowledge base about both is limited.

How Much Do Cyber Attacks Cost The Victims?

In the MafiaBoy denial of service attack on eBay, Yahoo, Dell and several other marquee web sites in February of 2000, each victim confidentially reported its actual losses to the FBI. I know a little about that case because I was the expert witness for the prosecution in MafiaBoy's trial. The technical attack on each victim was basically identical, and the outages were roughly the same length, but victims reported radically different estimates of damage. Their estimates ranged from zero to $5,000,000 depending on whether they included lost revenue, damage to reputation, management time, the direct costs of staff involved in the investigation and recovery, or none of those. Estimating losses is much more of an art than a science.

Another example of the difficulty of estimating losses was illustrated by the Nimda worm that raged through the Internet seven days, nearly to the minute, after the first airliner crashed into the World Trade Center. I interviewed more than a dozen victims confidentially, and they consistently told me the damage they incurred was between $300 and $700 per system – the actual cost of removing the infections from the systems and reinstalling software and data. For 150,000 infected systems, that adds up to about $75 million dollars. Yet within days of the attack, an economics firm was telling the press that the price tag was $835 million. Other people gave

estimates exceeding $2 billion. Before policy makers can rely on any damage assessments, a common protocol for damage estimation is needed. DHS can help develop that protocol.

How Important Are Simulations and Exercises?

Simulations and exercises are both valuable for improving America's effectiveness in responding to cyber attacks. The mathematical models simulating worms, created by organizations like CAIDA (Cooperative Association for Internet Data Analysis) at the University of California's San Diego Supercomputer Center, were instrumental in giving policy makers effective projections of the numbers of systems that would ultimately be infected by various worms. That kind of knowledge is extraordinarily valuable in the pressure cooker atmosphere of a worm infestation.

Simulating attacks through real world exercises are just as important for two reasons. The first reason is that emergency response systems rarely work as they were designed to operate. A few months ago a past deputy director of the House Information Systems (now called House Information Resource) told me a story about an exercise testing their fire emergency response plans. He wanted to ensure his organization would respond appropriately if a fire broke out in the building, so he scheduled a fire drill. When the alarm went off, most people, following patterns most of us developed in grade school, left, but no one in the computer room reacted at all. In a real fire they would probably have died. The problem: for some reason, in wiring the computer room, the electricians disconnected the power to the horns that sounded alarms. The computer room staff never heard the alarm. Without an exercise, no one would have known.

The other reason to run exercises involves the cyber dimension of physical attacks. Recall that in the aftermath of the September 11 attack, not only buildings were destroyed. The networks and systems of the New York Stock Exchange and all of Wall Street were a shambles. Without rapid reconstitution, the negative economic impact of the September 11 attack would have been even greater than it was. Verizon staff worked with the city's leaders 24 hours a day every day to rebuild the telephone and cyber networks needed to get trading restarted on Wall Street. They did an extraordinary job under difficult conditions, and a substantial part of their success was made possible because Verizon had already built a strong relationship with the mayor's office and the emergency response teams through planning and exercising disaster recovery protocols. Most cyber teams have no such connection with first responders and they need to know one another before an incident occurs..

We cannot have those groups exchanging business cards after an attack. The first responders will do a better job of planning if they know the cyber experts who can help them recover their networks and the issues those people will face when responding to emergencies. At the same time, the cyber people will be better team members if they understand what the mayors and governors need and jointly develop the action plans. We need to give the cyber people a seat at the table when planning for emergencies. These groups should learn from each other in advance, test communication paths and their ability to work together, identify problems and potential solutions, learn how long things take and how to speed them up. Exercises are the best way to make that happen.

What Can DHS Do?

Finally, you asked about how the Department of Homeland Security should work with the private sector in improving response and recovery. Let me summarize two key suggestions that go beyond the recommendations I covered earlier:

1. A central goal of the Department's cyber initiative should be to provide a single, technically savvy coordination point that the experts can rally around in responding to major attacks. I have been extremely impressed by the quality of people the Department has recruited to set up and run the new Cyber Security Tracking, Analysis, & Response Center (CSTARC). That group proved it can do extraordinary work in bringing together the public and private sector, both in responding to a vulnerability in sendmail and in the Slammer worm response. The key to CSTARC's long-term success is establishing a core group of very skilled people who then build a network of experts inside and outside government. Through exercises and responding to actual attacks, this community of cyber first-responders can create protocols and allocate responsibility for isolating malicious code, analyzing it, developing automated diagnostic and repair tools, and disseminating the tools and knowledge to the right people very rapidly.

2. As important as response and recovery are, prevention should have equal priority. DHS should allocate a large share of its time, attention, and budget to reducing the cyber vulnerabilities this nation faces. DHS can help by encouraging and supporting development of consensus benchmarks for safer configurations, but the Department's greatest impact on vulnerability reduction will come from persuading vendors of software, hardware, and network services that the government is serious about buying and running safer systems. The federal government is the only buyer large enough to get the attention of big vendors. DHS should make it clear, through both talk and action, that success in selling to the federal government is contingent upon delivering safely configured systems and automating the process of keeping those systems secure over time.

Thank you again for inviting me today and for your leadership in holding these hearings. I would be happy to try to answer any questions you might have.